

## Hierarchical wavelet modelling of environmental sensor data

Yann Ruffieux and A. C. Davison

*Ecole Polytechnique Fédérale de Lausanne*

**Abstract.** Motivated by the need to smooth and to summarize multiple simultaneous time series arising from networks of environmental monitors, we propose a hierarchical wavelet model for which estimation of hyperparameters can be performed by marginal maximum likelihood. The result is an empirical Bayes thresholding procedure whose results improve on those of wavethresh in terms of mean square error. We apply the approach to data from the SensorScope environmental modelling system, and briefly discuss issues that arise concerning variance estimation in this context.

### 1 Introduction

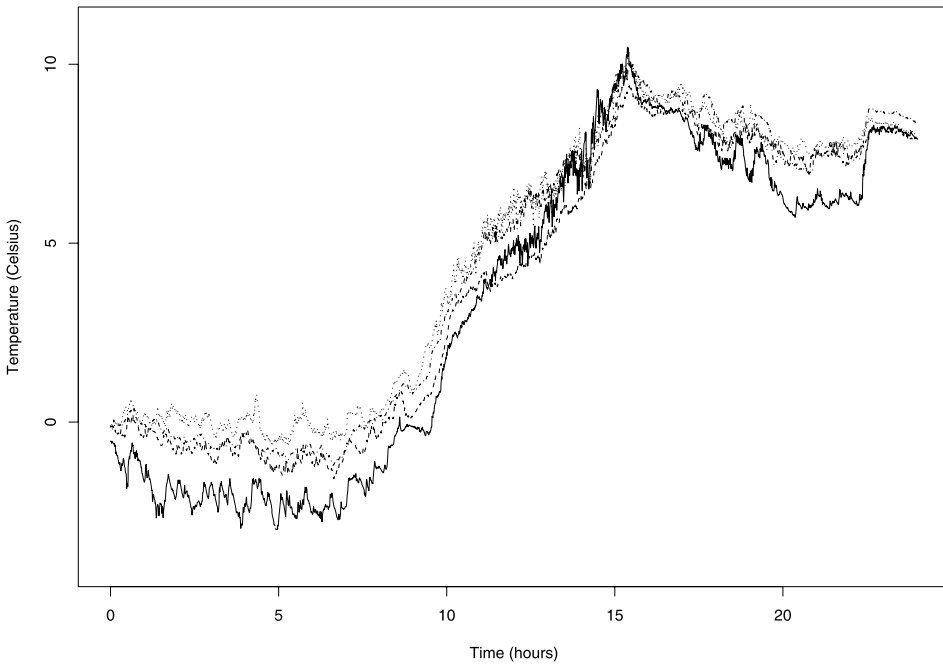
Realistic environmental modelling depends critically on detailed data that have hitherto largely been too expensive to obtain. For example, in temperature or rain-fall studies it has not been uncommon to base modelling on time series of daily values gathered many kilometres apart, so that interpolation, or downscaling, has been required to make inferences at more local scales, both in time and in space. The stunning decrease in the cost and increase in the quality of telecommunications and other electronic equipment now make it feasible to create dense wireless networks of cheap sensors that yield measurements at high spatial and temporal resolutions, with the potential to learn vastly more about the detailed working of environmental phenomena. SensorScope (<http://www.sensorscope.ch>) is an interdisciplinary project begun at EPFL which provides such data, the quantity and quality of which create new problems for statistics: of design of such networks; of data treatment; and of interpretation. Below we describe a simple approach to summarization of parallel data from many sensors, preliminary to more detailed exploitation.

Before deployment in more demanding surroundings, in July 2006 a working group embarked on a project to measure numerous atmospheric variables around the campus. The Lausanne Urban Canopy Experiment (LUCE) consisted of a network of around 100 weather stations, deployed within an area of roughly one-half of a square kilometre. Stations with various different configurations were used, collecting data on quantities such as air temperature, ground temperature, soil moisture, humidity, rainfall, and wind speed and direction, in real time and at short time

---

*Key words and phrases.* Empirical Bayes, environmental sensor, hierarchical model, mixture model, normal distribution, SensorScope, spike-and-slab model, wavelet.

Received May 2011; accepted May 2011.



**Figure 1** Twenty-four hours of air temperature measurements ( $^{\circ}\text{C}$ ) at four SensorScope weather stations on the EPFL campus.

intervals. The resulting data are noteworthy: they are highly localized both in time and space; they run over several months; and they are gathered in an unusual urban-type environment. Similar dense networks have been or are being deployed in more taxing settings, for example, at the Grand-St-Bernard pass (altitude 2400 m) and at the Plaine Morte glacier (altitude 2750 m); more details can be obtained from the link given above.

Figure 1 displays ambient temperatures taken during the LUCE experiment at four stations in a time frame of 24 hours (midnight to midnight); the measurements were made roughly every 30 seconds. The curves show a common trend due to the high proximity of the stations but have distinct features due to their particular surroundings—exposure, altitude, ground type, and so forth. The signals have appreciable local variation and some kind of smoothing seems required.

Although smoothing the signals independently would be quick and easy, we would like to allow for the obvious similarity between signals from different stations. One approach is to assume the existence of an underlying curve that generates the station-specific curves. We will concentrate on the air temperature measurements, as a prototype for other series of continuous values. We define temperature curves  $Y_s(t)$  for each of the  $S$  weather stations, and assume that

$$Y_s(t) = \mu(t) + \psi_s(t) + e_s(t), \quad s = 1, \dots, S.$$

The time variable  $t$  can be taken over any interval, but for our purposes it covers 24 hours, from midnight to midnight. The function  $\mu(t)$  can be regarded as the underlying temperature curve for the whole campus and  $\psi_s(t)$  as the effect from station  $s$ , while  $e_s(t)$  represents the residual effect function at station  $s$ . Thus the function  $\mu(t) + \psi_s(t)$  can be considered to be the noise-free curve for station  $s$ .

In practice data may be observed at irregular intervals, at times varying slightly from station to station. For later developments we interpolate the observed time series to obtain observations on a regular grid  $t_1, \dots, t_N$  over the whole day, with  $t_{i+1} - t_i = c$  for all  $i$  and with  $N = 2^J$  for a fixed positive integer  $J$ . The model may then be written

$$Y_s = \mu + \psi_s + e_s, \quad s = 1, \dots, S, \quad (1.1)$$

with  $Y_s = [Y_s(t_1), \dots, Y_s(t_N)]^T$  denoting the vector of observations at station  $s$ ,  $\mu = [\mu(t_1), \dots, \mu(t_N)]^T$ ,  $\psi_s = [\psi_s(t_1), \dots, \psi_s(t_N)]^T$ , and  $e_s = [e_s(t_1), \dots, e_s(t_N)]^T$ . Our goal here will be to filter out the noise  $e_s$  in order to find smooth estimators for  $\mu$  and  $\psi_s$ .

Equations (1.1) can be written in matrix form as

$$Y = XB + E, \quad (1.2)$$

where  $Y$  is the  $S \times N$  matrix with the observation vectors  $Y_s$  in rows,  $X = [1_S | I_S]$  is a  $S \times (S + 1)$  design matrix,  $B$  is a  $(S + 1) \times N$  matrix with  $\mu$  in the first row and  $\psi_s$  in row  $s + 1$ , and  $E$  is the  $S \times N$  matrix of residuals. Here and elsewhere, the term  $1_S$  designates a  $S \times 1$  vector of ones and  $I_S$  is the  $S \times S$  identity matrix. Expression (1.2) corresponds to a version of the functional mixed-effects model described in [Morris and Carroll \(2006\)](#), but with no random effect. A time-domain approach to such problems using kernel density estimation has been described by [Ghosh \(2001\)](#).

One of our aims is to develop an automatic approach: ultimately we would like to produce the necessary estimates for SensorScope in real time without any intervention. Many of the assumptions below are made towards that end. In Section 2 we outline our model and then in Section 3 we discuss inference for its hyperparameters. Section 4 describes a small simulation study to compare our approach with use of `wavethresh`, and is followed by an application to the SensorScope data. The paper concludes with a brief discussion.

## 2 Wavelet regression and Bayesian modelling

We work in the wavelet domain, which gives an attractive basis for curve regularization and modelling. Wavelets have received a lot of attention from the mathematical and statistical communities in the past few years, and much has been written about them. [Strang \(1993\)](#) gives a nice introduction, while [Percival and Walden \(2000\)](#) provide a detailed account focusing on time series analysis.

Let  $W$  be the  $N \times N$  orthogonal discrete wavelet transform matrix under a given wavelet basis. Then right-multiplication of a  $1 \times N$  vector by  $W^T$  corresponds to a change of basis from the time domain to the wavelet domain. The coefficients in this new basis have a specific interpretation in terms of local variation of the signal, much like Fourier coefficients. However the wavelet basis allows the local variation to depend on location, or, in our context, time. There is an infinite number of possible wavelet bases, but only a few of them are widely used (Daubechies, 1992). At this point no assumption is needed on the type of wavelet used, and  $W$  can be seen as a generic change of basis. Wavelet coefficients can be computed in a highly efficient manner using the pyramid algorithm (Mallat, 1989).

Right-multiplying both sides of (1.2) by  $W^T$  yields

$$D = XB^* + E^*, \tag{2.1}$$

where  $D = YW^T$  is a  $S \times N$  matrix whose row  $s$  contains the ‘‘observed’’ wavelet coefficients for station  $s$ ,  $B^* = BW^T$  contains the wavelet coefficients for the mean function  $\mu$  in the first row and the wavelet coefficients for the station-specific effects  $\psi_s$  in the other rows, while  $E^* = EW^T$  contains the residuals in the wavelet space. We double-index the rows of  $D$ ,  $B^*$ , and  $E^*$  to include the scale and location of each coefficient: for  $j = 0, 1, \dots, J - 1$ ,  $k = 1, \dots, 2^j$  and  $s = 1, \dots, S$ , write  $d_{jk}^{(s)}$  as the observed coefficient at scale  $j$  and location  $k$  in the wavelet decomposition of  $Y_s$ . Then a coefficient with large  $j$  has high resolution and so is more likely to be interpreted as noise. Similarly, define  $\theta_{jk}^{(s)}$  and  $\varepsilon_{jk}^{(s)}$  as the  $(j, k)$ -coefficient in the wavelet decompositions of  $\psi_s$  and  $e_s$  respectively. Finally set  $\zeta_{jk}$  as the  $(j, k)$ -coefficient in the decomposition of the mean vector  $\mu$ . For fixed  $j, k$ , and  $s$ , the relation between these coefficients is

$$d_{jk}^{(s)} = \zeta_{jk} + \theta_{jk}^{(s)} + \varepsilon_{jk}^{(s)}; \tag{2.2}$$

we assume that the  $\varepsilon_{jk}^{(s)}$  are mutually independent, identically distributed variables from a centred Gaussian density of variance  $\sigma^2$ . There is also a scaling coefficient for each decomposition in expression (2.2); see below.

We now set priors on the parameters  $\zeta_{jk}$  and  $\theta_{jk}^{(s)}$ . We first define a Bernoulli random variable  $Z_{jk}$ , with  $P(Z_{jk} = 1) = \pi_j$  and  $P(Z_{jk} = 0) = 1 - \pi_j$ , and link it to the coefficients as follows:

- if  $Z_{jk} = 1$ , then we independently set  $\zeta_{jk} \sim \mathcal{N}(0, \tau_j^2)$  and  $\theta_{jk}^{(s)} \sim \mathcal{N}(0, \eta^2)$  for  $s = 1, \dots, S$ ;
- if  $Z_{jk} = 0$ , then  $\zeta_{jk}$  and each  $\theta_{jk}^{(s)}$  follow a degenerate distribution with unit mass at zero.

This latent variable  $Z_{jk}$  thus indicates whether the corresponding coefficients  $\zeta_{jk}$  and  $\{\theta_{jk}^{(s)}, s = 1, \dots, S\}$  are ‘‘switched on.’’ We use microarray terminology and say that the  $(j, k)$ -coefficients are differentially expressed if  $Z_{jk} = 1$ . The main

purpose in making the above assumptions is that we consider that only those coefficients which are differentially expressed contain information on the smooth signal, the rest being treated as noise. In particular, we expect  $\pi_j$  to decrease with  $j$  so that high-resolution coefficients are less likely to be differentially expressed *a priori*.

On marginalizing over  $Z_{jk}$  we see that the coefficients follow mixture distributions, that is,

$$\zeta_{jk} \sim \pi_j \mathcal{N}(0, \tau_j^2) + (1 - \pi_j)\delta_0$$

and

$$\theta_{jk}^{(s)} \sim \pi_j \mathcal{N}(0, \eta^2) + (1 - \pi_j)\delta_0,$$

with  $\delta_0$  denoting the distribution with unit mass at zero. These prior assumptions are very similar to those of [Morris and Carroll \(2006\)](#), who also set a mixture model on the elements of  $B$ . An important difference here is that for fixed  $(j, k)$ , the coefficients  $\zeta_{jk}$  and  $\theta_{jk}^{(s)}$ ,  $s = 1, \dots, S$  are marginally dependent: *a priori* they are either all differentially expressed or all zero. This mixture approach was also used by [Abramovich, Sapatinas and Silverman \(1998\)](#) and [Johnstone and Silverman \(2005\)](#), but applied to single times series.

For completeness, we place a vague prior on the single scaling coefficient of each wavelet decomposition. Thus these coefficients are estimated *a posteriori* by the sample scaling coefficients, that is, the sample mean multiplied by the square root of the number of coefficients.

Now we establish the posterior distribution of the wavelet coefficients given the observed coefficients  $D$  and the hyperparameters  $\sigma^2, \eta^2, \tau_j^2$ , and  $\pi_j$ . Set  $d_{jk} = [d_{jk}^{(1)}, \dots, d_{jk}^{(S)}]^T$ , a column of the matrix  $D$ , and  $\theta_{jk} = [\theta_{jk}^{(1)}, \dots, \theta_{jk}^{(S)}]^T$ . It is easily seen that, given  $Z_{jk} = 1$ , the joint distribution of  $[d_{jk}^T, \theta_{jk}^T, \zeta_{jk}]^T$  is a centred multivariate Gaussian distribution, and that it is degenerate given  $Z_{jk} = 0$ . In particular, the marginal density function of  $d_{jk}$  is the mixture

$$p(d_{jk}) = \pi_j g(d_{jk}; 0, A_j) + (1 - \pi_j)g(d_{jk}; 0, \sigma^2 I_S), \tag{2.3}$$

where  $g(\cdot; m, \Sigma)$  is the density function of a multivariate Gaussian random variable with mean  $m$  and covariance matrix  $\Sigma$ , and where

$$A_j = (\sigma^2 + \eta^2)I_S + \tau_j^2 1_S 1_S^T$$

is an equicorrelation matrix, whose inverse and determinant are  $(\sigma^2 + \eta^2)^{-1}I_S - \tau_j^2 / \{(\sigma^2 + \eta^2)(\sigma^2 + \eta^2 + S\tau_j^2)\} 1_S 1_S^T$  and  $(\sigma^2 + \eta^2)^{2(S-1)}(\sigma^2 + \eta^2 + S\tau_j^2)$ , respectively. It is straightforward to find the posterior odds  $\omega_{jk}$  of the  $(j, k)$ -coefficients being differentially expressed:

$$\omega_{jk} = \frac{P(Z_{jk} = 0|d_{jk})}{P(Z_{jk} = 1|d_{jk})} = \frac{1 - \pi_j}{\pi_j} \times \frac{g(d_{jk}; 0, \sigma^2 I_S)}{g(d_{jk}; 0, A_j)}.$$

Using properties of the multivariate normal distribution, we find that the posterior distribution function of  $\zeta_{jk}$  is the mixture

$$F(\zeta_{jk}|d_{jk}) = \frac{1}{1 + \omega_{jk}} \Phi \left\{ \frac{\zeta_{jk} - \tau_j^2 \mathbf{1}_S^T A_j^{-1} d_{jk}}{\tau_j \sqrt{1 - \tau_j^2 \mathbf{1}_S^T A_j^{-1} \mathbf{1}_S}} \right\} + \frac{\omega_{jk}}{1 + \omega_{jk}} I\{\zeta_{jk} \geq 0\}, \quad (2.4)$$

where  $\Phi$  is the standard normal distribution function and  $I$  is the indicator function. Recall that for fixed  $s$ , the  $\theta_{jk}^{(s)}$  represent the coefficients for the *effect* of station  $s$ . We will focus on the estimation of  $v_{jk}^{(s)} = \zeta_{jk} + \theta_{jk}^{(s)}$ : for fixed  $s$ , the  $v_{jk}^{(s)}$  are the wavelet coefficients for the noise-free signal from station  $s$ . Let  $c_{js} = [\tau_j^2, \dots, \tau_j^2, \tau_j^2 + \eta^2, \tau_j^2, \dots, \tau_j^2]^T$  be a  $S \times 1$  vector, with the  $\tau_j^2 + \eta^2$  term in position  $s$ . The posterior distribution of  $v_{jk}^{(s)}$  is

$$F(v_{jk}^{(s)}|d_{jk}) = \frac{1}{1 + \omega_{jk}} \Phi \left\{ \frac{v_{jk}^{(s)} - c_{js}^T A_j^{-1} d_{jk}}{\sqrt{\eta^2 + \tau_j^2 - c_{js}^T A_j^{-1} c_{js}}} \right\} + \frac{\omega_{jk}}{1 + \omega_{jk}} I\{v_{jk}^{(s)} \geq 0\}. \quad (2.5)$$

These distribution functions have jumps of sizes  $\omega_{jk}/(1 + \omega_{jk})$  at  $v_{jk}^{(s)} = 0$ .

### 3 Hyperparameter estimation and inference

#### 3.1 Basic model

We use an empirical Bayes approach to choose the hyperparameters  $\sigma^2, \eta^2, \tau_j^2$ , and  $\pi_j$  of our model. The full marginal log likelihood for these hyperparameters is

$$\ell(\sigma^2, \eta^2, \tau_0^2, \dots, \tau_{J-1}^2, \pi_0, \dots, \pi_{J-1}; D) = \sum_{j,k} \log p(d_{jk}), \quad (3.1)$$

where the  $p(d_{jk})$  are computed from (2.3). The hyperparameter estimates will be the values which maximize (3.1). The number of hyperparameters,  $2(J + 1)$ , can be relatively high, since typically  $J \geq 10$  in this context. We can ease the computation somewhat by following the model of Abramovich, Sapatinas and Silverman (1998) for the variance of  $\zeta_{jk}$ ,

$$\tau_j^2 = 2^{-j\alpha} C, \quad j = 0, 1, \dots, J - 1, \quad (3.2)$$

for  $\alpha, C > 0$ . Making the appropriate substitutions in (3.1), we can then maximize over the parameters  $\alpha$  and  $C$  rather than over  $\tau_0^2, \dots, \tau_{J-1}^2$ . Experiments with unconnected  $\tau_j^2$  suggest that (3.2) is reasonable. We expect only the higher-resolution coefficients to correspond to noise, so we group the  $k$  lower scale  $\pi_j$ 's into a single parameter  $\pi_B$ , which we expect to be close to unity.

To estimate the wavelet coefficients  $\zeta_{jk}$  and  $v_{jk}^{(s)}$ , we follow the Abramovich, Sapatinas and Silverman (1998) thresholding approach, by computing the posterior

medians of the distributions (2.4) and (2.5). We can find the median for the mixture distribution  $H(x) = (1 + \omega)^{-1}\Phi\{(x - \mu)/\nu\} + \omega/(1 + \omega)I(x \geq 0)$  corresponding to (2.4) and (2.5) as follows:

- if  $\omega \geq 1$ , then the jump at 0 is greater than 1/2 and the median is zero;
- if  $(1 - \omega)/2 \leq \Phi(-\mu/\nu) \leq (1 + \omega)/2$ , the jump at 0 starts below 1/2 and lands above 1/2, so the median is zero;
- otherwise the median is

$$\mu + \nu\Phi^{-1}\left[\frac{1}{2} + \frac{\omega}{2}\text{sign}\left\{\frac{\Phi(-\mu/\nu)}{1 + \omega} - \frac{1}{2}\right\}\right].$$

Let  $\bar{d}_{jk} = \sum_{s=1}^S d_{jk}^{(s)}/S$  be the mean of the  $(j, k)$ -coefficients across the sites. After some algebra we find that

$$\hat{\zeta}_{jk} = \text{med}(\zeta_{jk}|d_{jk}) = \text{sign}(\bar{d}_{jk}) \max(0, \gamma_{jk}),$$

with

$$\gamma_{jk} = \frac{S\tau_j^2}{\sigma^2 + \eta^2 + S\tau_j^2}|\bar{d}_{jk}| - \sqrt{\frac{\tau_j^2(\sigma^2 + \eta^2)}{\sigma^2 + \eta^2 + S\tau_j^2}} \times \Phi^{-1}\left\{\frac{1 + \min(\omega_{jk}, 1)}{2}\right\}.$$

The “threshold” rule is thus

$$\hat{\zeta}_{jk} = 0 \iff |\bar{d}_{jk}| \leq \frac{1}{S}\sqrt{\frac{(\sigma^2 + \eta^2 + S\tau_j^2)(\sigma^2 + \eta^2)}{\tau_j^2}} \times \Phi^{-1}\left\{\frac{1 + \min(\omega_{jk}, 1)}{2}\right\}.$$

As  $\omega_{jk}$  depends both on  $\bar{d}_{jk}$  and on  $\sum_{s=1}^S (d_{jk}^{(s)})^2$ , this is not a pure threshold rule.

As a result of taking the posterior median, the estimated coefficients are either differentially expressed or set to zero, as in the prior model. By setting certain coefficients to zero we remove the noise and retain just the smooth components of the curves.

Once we have our estimates  $\hat{\zeta}_{jk}$  and  $\hat{\nu}_{jk}^{(s)}$ , we reconstruct the smooth signal by applying the inverse discrete wavelet transform to the appropriate vectors of estimated coefficients: the smooth estimated curves  $\hat{\mu}$  and  $\hat{\psi}_s$  can be extracted from the matrix  $W\hat{B}^*$ , where  $\hat{B}^*$  is the matrix  $B^*$  with the estimates  $\hat{\zeta}_{jk}$  and  $\hat{\theta}_{jk}^{(s)} = \hat{\zeta}_{jk} - \hat{\nu}_{jk}^{(s)}$  inserted at the appropriate locations.

### 3.2 Nonconstant noise variance

Above we assumed a constant variance for the noise. However examination of the daily temperature data in both time and wavelet domains suggests that there are two variance regimes: one towards the afternoon during which the variance is

clearly larger, and another with less variability. The boundaries between them are hard to pinpoint, and vary from day to day. We deal with this by assigning mixtures of normal densities to the errors  $\varepsilon_{jk}^{(s)}$ , namely:

$$\varepsilon_{jk}^{(s)} \sim \tilde{\pi} \mathcal{N}(0, \sigma_1^2) + (1 - \tilde{\pi}) \mathcal{N}(0, \sigma_2^2),$$

with  $0 < \tilde{\pi} < 1$ . We have implicitly defined a Bernoulli variable  $\tilde{Z}_{jk}$  such that  $\tilde{Z}_{jk} = 1$  implies that the error variance equals  $\sigma_1^2$  while  $\tilde{Z}_{jk} = 0$  implies it equals  $\sigma_2^2$ . As in the constant-variance case,  $\sigma_1^2$  and  $\sigma_2^2$  can be estimated by maximizing the marginal likelihood

$$\ell(\sigma_1^2, \sigma_2^2, \eta^2, \tau_j^2, \pi_j, \tilde{\pi}; D) = \sum_{j,k} \log p(d_{jk}),$$

with

$$p(d_{jk}) = \pi_j \tilde{\pi} g(d_{jk}; 0, A_{1j}) + \pi_j (1 - \tilde{\pi}) g(d_{jk}; 0, A_{2j}) \\ + (1 - \pi_j) \tilde{\pi} g(d_{jk}; 0, \sigma_1^2 I_S) + (1 - \pi_j) (1 - \tilde{\pi}) g(d_{jk}; 0, \sigma_2^2 I_S),$$

where  $A_{1j}$  is the matrix  $A_j$  with the  $\sigma^2$  terms replaced by  $\sigma_1^2$ , and similarly for  $A_{2j}$  and  $\sigma_2^2$ . The posteriors for  $\zeta_{jk}$  and  $v_{jk}^{(s)}$  are now mixtures of two distinct normal components and a point mass at zero, making computation of the posterior medians awkward. We circumvent this by first computing the posterior odds of the error variance being  $\sigma_1^2$ , for each  $(j, k)$ , giving

$$\tilde{\omega}_{jk} = \frac{1 - \tilde{\pi}}{\tilde{\pi}} \times \frac{\pi_j g(d_{jk}; 0, A_{2j}) + (1 - \pi_j) g(d_{jk}; 0, \sigma_2^2 I_S)}{\pi_j g(d_{jk}; 0, A_{1j}) + (1 - \pi_j) g(d_{jk}; 0, \sigma_1^2 I_S)}.$$

The estimators are then

$$\hat{\zeta}_{jk} = \text{med}(\zeta_{jk} | \tilde{Z}_{jk} = 1) I\{\tilde{\omega}_{jk} \leq 1\} + \text{med}(\zeta_{jk} | \tilde{Z}_{jk} = 0) I\{\tilde{\omega}_{jk} > 1\}$$

and

$$\hat{v}_{jk}^{(s)} = \text{med}(v_{jk}^{(s)} | \tilde{Z}_{jk} = 1) I\{\tilde{\omega}_{jk} \leq 1\} + \text{med}(v_{jk}^{(s)} | \tilde{Z}_{jk} = 0) I\{\tilde{\omega}_{jk} > 1\}.$$

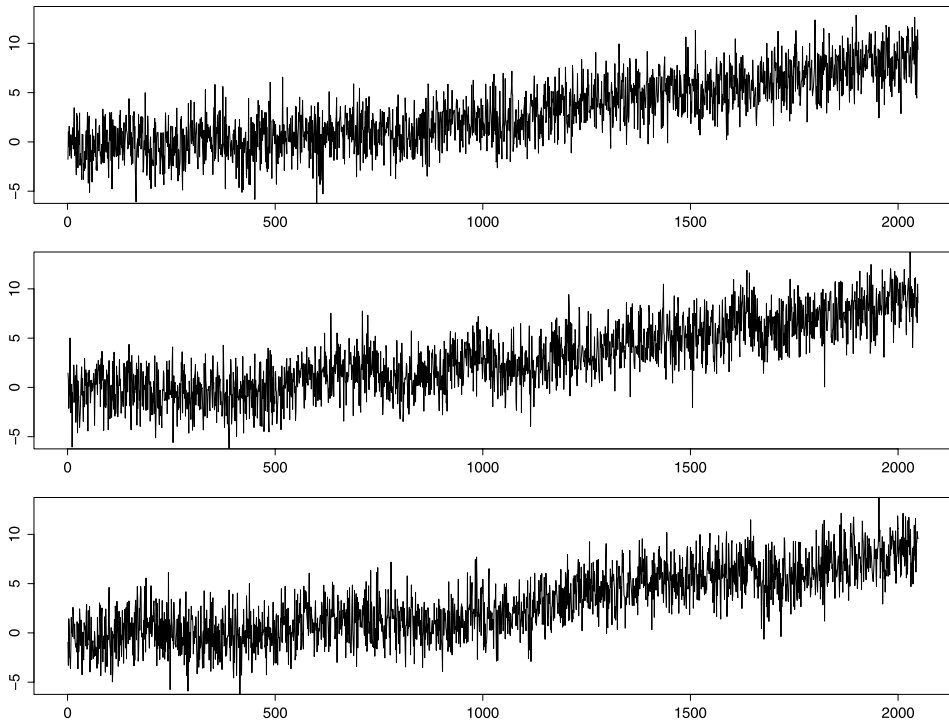
This dual-variance model need not be applied at all resolutions, but only at the few highest which directly correspond to noise. One can use the single-variance model (2.3) for the lower-scale levels.

### 4 Simulation study

Here we apply our estimation approach to simulated data and compare the results with independent wavelet thresholding.

We created  $S = 8$  sets of  $2^J$  wavelet coefficients following the model described in Section 3, with  $J = 11$ . More precisely, we generated “true” coefficients  $\zeta_{jk}$

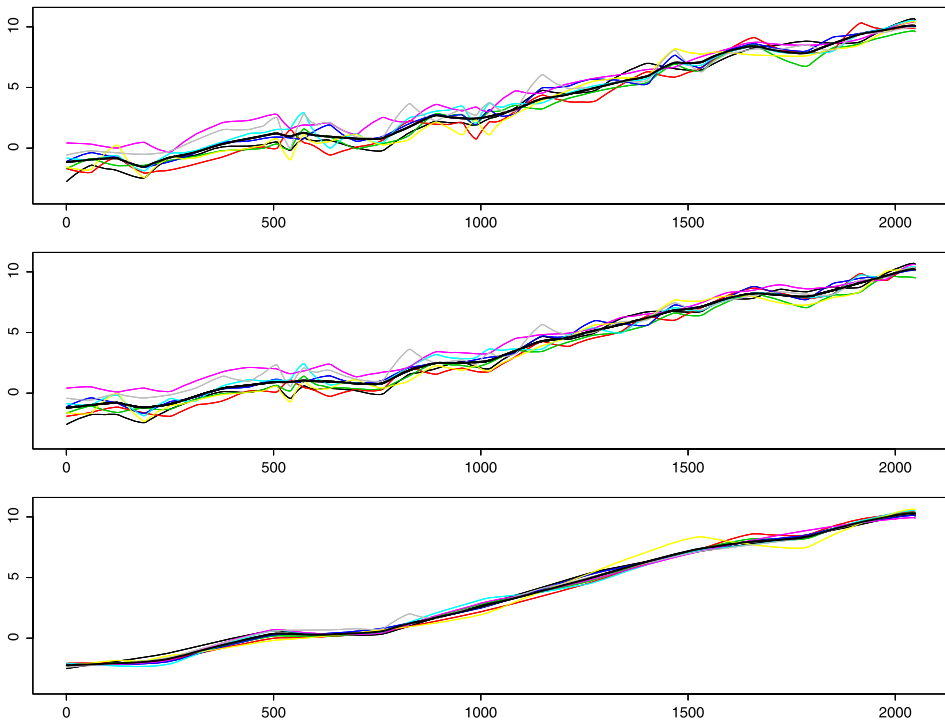




**Figure 2** Three time series obtained by applying the inverse discrete wavelet transform to sets of simulated wavelet coefficients.

and  $v_{jk}^{(s)}$  from the appropriate mixture models and added Gaussian noise to obtain the  $d_{jk}^{(s)}$ . The coefficient scale parameter used was that of the coefficients from the SensorScope temperature measurements. Setting  $\sigma^2 = 4$ ,  $\eta^2 = 15$ ,  $C = 1000$ , and  $\alpha = 2$ , with  $\pi_j$  quickly decreasing in  $j$ , the reconstructed simulated time series, while very noisy, have a similar underlying smooth curve; three are shown in Figure 2. Applying empirical Bayes estimation yields  $\hat{\sigma}^2 = 3.935$ ,  $\hat{\eta}^2 = 12.15$ ,  $\hat{C} = 1162$ , and  $\hat{\alpha} = 2.729$ . The fairly large error in the latter three estimates can probably be explained by the fact that they are based only on a few low-scale, differentially expressed coefficients, whereas the noise variance estimate is based on numerous high-resolution coefficients. Further simulations based on the above “true” hyperparameter values suggest that, apart from  $\hat{\eta}^2$ , the estimates are unbiased. With the values from which we simulate, the empirical Bayes approach seems to underestimate  $\eta^2$ .

Figure 3 compares the “true” curves underlying the data in Figure 2, obtained by applying the inverse discrete wavelet transform to the simulated coefficients  $\zeta_{jk}$  and  $v_{jk}^{(s)}$ , with the corresponding estimates reconstructed from the posterior medians of (2.4) and (2.5), using the empirical Bayes hyperparameter values. None

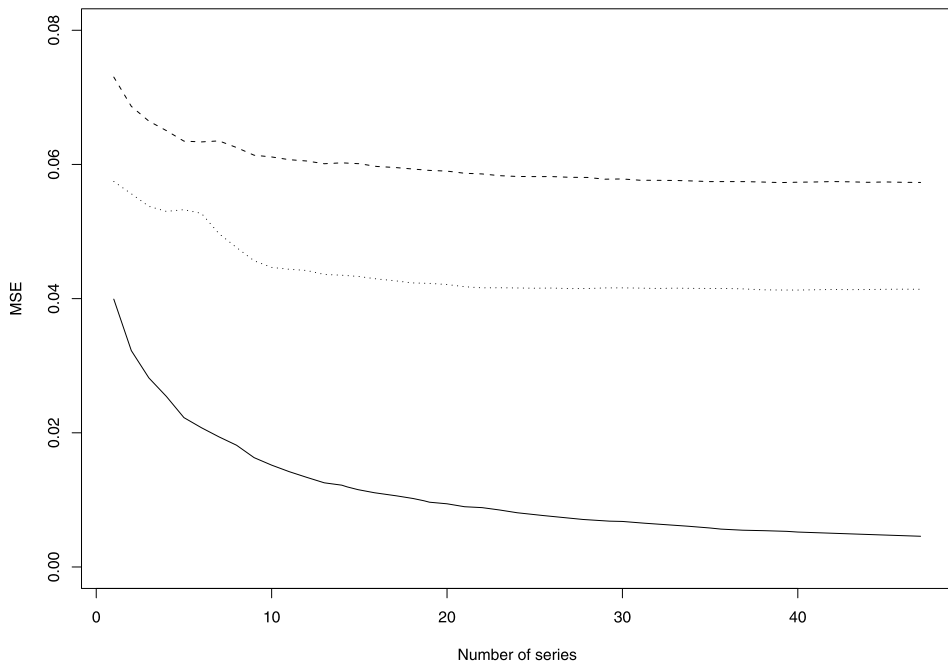


**Figure 3** Reconstructed time series obtained by applying inverse discrete wavelet transform to, from top to bottom, the true wavelet coefficients, the coefficients estimated with the median rule via empirical Bayes, and the coefficients after thresholding with `wavethresh` using default settings. Narrow lines: the eight station-specific smooth signals, thick line: the underlying mean signal (in the bottom panel, the pointwise mean of the station-specific smooth signals).

of the estimates  $\hat{v}_{jk}^{(s)}$  is falsely declared to be differentially expressed or falsely declared to be zero. Of the  $2^{11}$  estimates for  $\zeta_{jk}$ , six are falsely declared to be zero, and one is falsely declared to be differentially expressed. Also shown are the curves smoothed independently using universal “hard” wavelet thresholding (Nason and Silverman, 1994); this clearly oversmooths, and thereby fails to capture most of the true features of the curves.

Given the nature of our model, we expect that a coefficient estimated from a given station may “borrow strength” from the estimates for the other stations. Thus having data from more stations should give better estimates, both of the underlying coefficients  $\zeta_{jk}$  and the station-specific coefficients  $v_{jk}^{(s)}$ . Figure 4 shows the corresponding mean square errors

$$\text{MSE}_\zeta(S) = \frac{1}{2^J} \sum_{j,k} (\zeta_{jk} - \hat{\zeta}_{jk,S})^2, \quad \text{MSE}_v(S) = \frac{1}{2^J S} \sum_{j,k} \sum_{s=1}^S (v_{jk}^{(s)} - \hat{v}_{jk,S}^{(s)})^2,$$

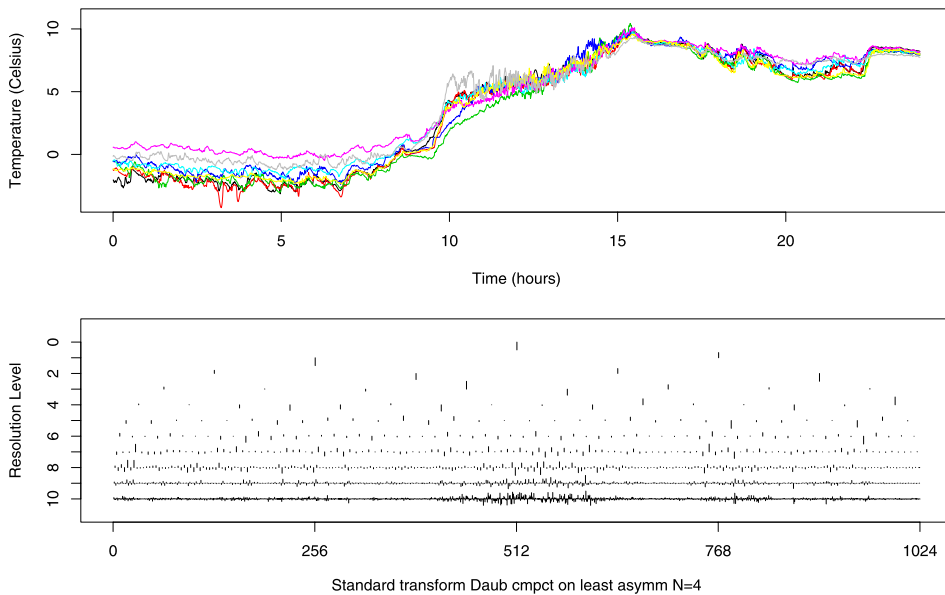


**Figure 4** Mean square errors (MSE) for different estimation methods, based on 50 replicates with different numbers of stations. The lines show: MSE of the posterior median estimates (solid) and wavethresh estimates (dashes) of the grand mean coefficients  $\zeta_{jk}$ , with the wavethresh estimates simply the mean of the thresholded coefficients; and MSE of the posterior median estimates of the signal-specific coefficients  $v_{jk}^{(s)}$  (dots). wavethresh estimates for these have an MSE of about 0.14, which does not decrease with the number of stations.

as functions of the number of stations  $S$ , where  $\hat{\zeta}_{jk,S}$  and  $\hat{v}_{jk,S}^{(s)}$  are the estimates based on  $S$  signals. As we increase  $S$ , we do not resimulate all the previous coefficients, but simply augment the data by adding further series of coefficients. The values for the hyperparameters were the same as above. We see that  $\text{MSE}_{\zeta}(S)$  decreases roughly exponentially. The evolution of the signal-specific estimates is less clear: the “borrowing of strength” seems to kick in at around  $S = 12$  before stabilizing. We would not see these improvements in performance if the signals were modelled independently.

## 5 Application to SensorScope measurements

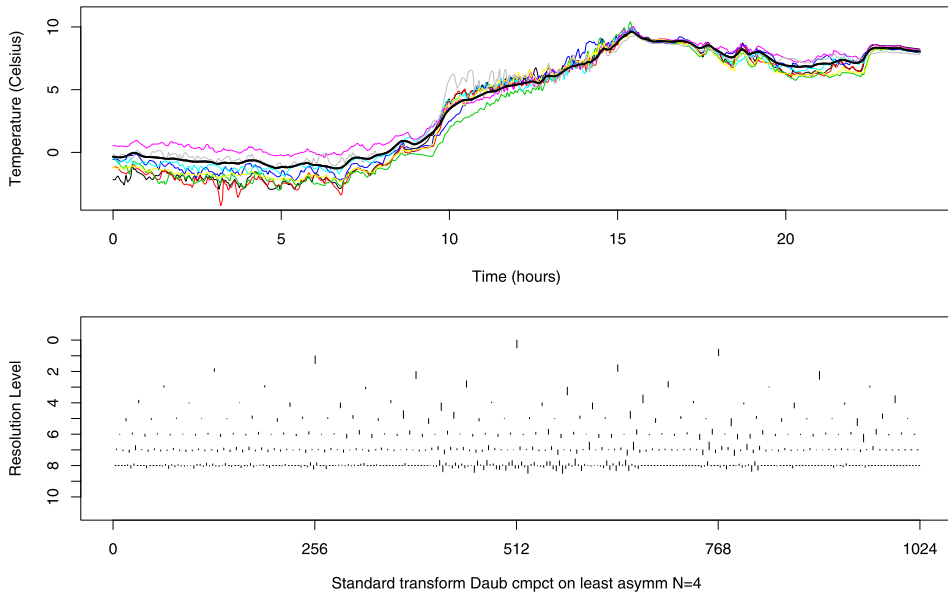
We now consider several ways of smoothing the SensorScope time series. The different approaches were applied to all  $S = 73$  series for which data were available on a particular day, though for clarity only eight will be plotted. We use Daubechie’s (1992) least-asymmetric compactly supported wavelet with four vanishing moments.



**Figure 5** Example series from the SensorScope data. Top: eight time series consisting of SensorScope air temperature data taken over 24 hours. Bottom: wavelet decomposition of one of the series.

Figure 5 shows the eight time series and a wavelet decomposition of one of them; note the diurnal fluctuations in the coefficient sizes at the lowest level. Figure 6 shows the result of applying the mixed variance model of Section 3.2 to the two finest coefficient levels. There is some smoothing of the individual original series, but little visual difference between this approach and the use of a single variance. The global curve estimate appears slightly biased because it is based on all 73 series, not merely on those plotted. The effect of the mixture is clearer in the wavelet domain: without the mixture, some coefficients at levels 9 and 10 survive thresholding, but with the mixture they are all zeroed out. As expected, the empirical Bayes estimates for the  $\pi_j$  decrease rapidly with the level  $j$ . We have merged the  $\pi_j$  for the five coarsest levels into a single parameter, whose estimate is practically unity; the estimates for  $\pi_9$  and  $\pi_{10}$  are both very close to zero. Moreover we find  $\hat{\eta}^2 = 0.98$ ,  $\hat{C} = 41761$ , and  $\hat{\alpha} = 3.15$ . As for the variance mixture parameters, the estimates of  $\sigma_a^2$ ,  $\sigma_b^2$ , and  $\tilde{\pi}$  are 0.0036, 0.028, and 0.75 respectively.

The top panel of Figure 7 displays the wavethresh smoothing of the eight time series using the default “hard” thresholding and error variance estimate based on the squared median absolute deviation of the coefficients at level 3 and higher. The results seem unsatisfactory in comparison with those of the hierarchical model: more detailed noise has survived thresholding. The middle panel shows the results of wavethresh when the error variance estimate is based on variance of the coefficients. The choice of variance estimate is crucial: with this much larger variance



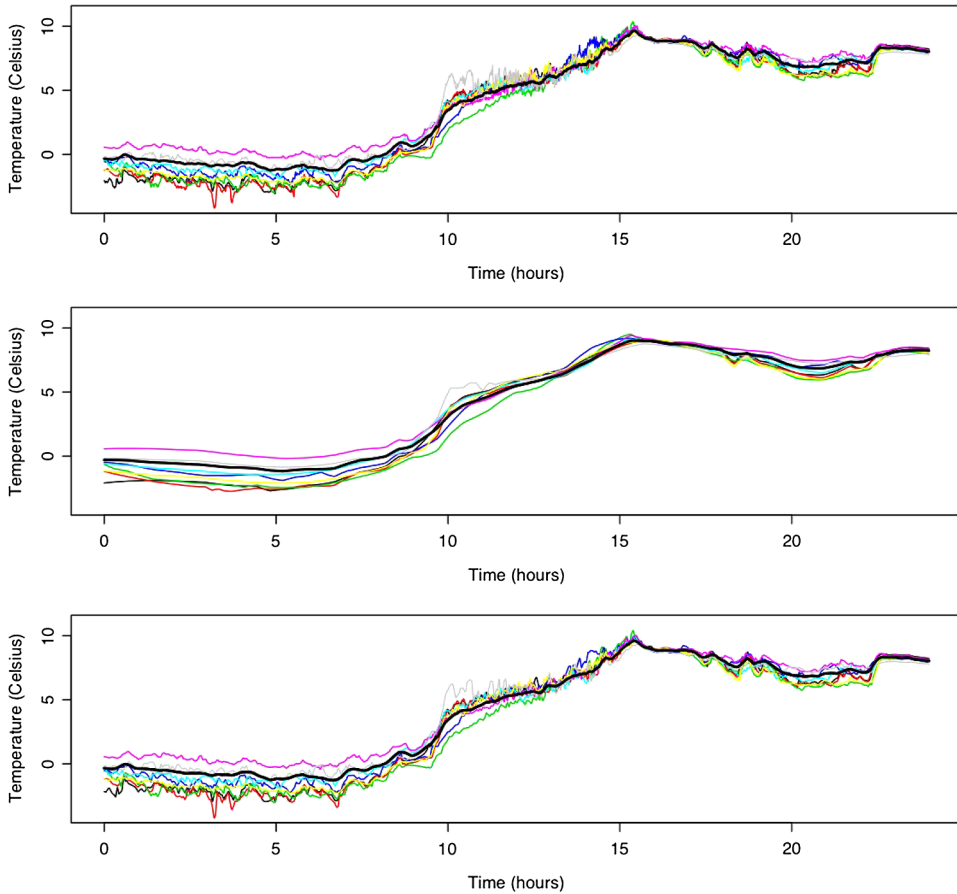
**Figure 6** Application of the hierarchical model to the SensorScope data. Top: reconstruction after applying the hierarchical model with a mixture for the error variance; the dark line represents the global curve; bottom: corresponding estimates of the wavelet coefficients  $\zeta_{jk}$ .

estimate the curves are appreciably smoother, with so few coefficients surviving that some individual wavelets can be identified. For some purposes these smoother curves might be preferred. The lower panel of the figure shows the results of our hierarchical approach with the noise variance fixed as the variance of all coefficients at levels 3 and higher; though not so smooth as in the second panel, the result might for some purposes be regarded as more satisfactory than the upper panel of Figure 6.

## 6 Discussion

We have described a simple hierarchical model for treating numerous mutually dependent time series. An empirical Bayes approach was used to assign values to the hyperparameters of the model, and the wavelet coefficients were estimated using a posterior median rule.

Experiments with simulated data show a clear gain in treating the time series simultaneously rather than independently. Any gain is less obvious when applying our methodology to the SensorScope data. The empirical Bayes estimate for the error variance is relatively small, so many high level coefficients pass the thresholding filter. A potential solution would be to assign a value to the error variance beforehand—for example, the variance of all the coefficients at level  $j_0$  and higher.



**Figure 7** *Effect of variance estimation on reconstruction. Top: wavethresh smoothing using default settings; middle: wavethresh smoothing with modified settings; bottom: reconstruction after applying the hierarchical model, with the error variance fixed beforehand as the variance of the coefficients at level 3 and higher.*

In this case  $j_0$  might be interpreted as a smoothing parameter, a “slider” that determines the strength of the smoothing.

One natural extension is to allow for a different variance at each station in the basic model (1.1). To test the effect of this we estimated the variances for each series using the median absolute deviation of its finest level of wavelet coefficients, rescaled the wavelet coefficients to have unit standard error, and applied the mixture model approach to the result, finally back-transforming to allow for the different variances. This had little effect on the reconstructed curves, however.

Other possibilities for work on the rich SensorScope database are the inclusion of covariates to allow for the particular surroundings of each station, and the treatment of outliers and missing data.

## Acknowledgments

We thank the Swiss National Science Foundation, the Francqui Foundation, and the ETH domain Centre for Competence in Environment and Sustainability EXTREMES for financial support, and an anonymous reader for comments on a previous version.

## References

- Abramovich, F., Sapatinas, T. and Silverman, B. W. (1998). Wavelet thresholding via a Bayesian approach. *Journal of the Royal Statistical Society Series B* **60**, 725–749. [MR1649547](#)
- Daubechies, I. (1992). *Ten Lectures on Wavelets*. Philadelphia: Society for Industrial and Applied Mathematics. [MR1162107](#)
- Ghosh, S. (2001). Nonparametric trend estimation in replicated time series. *Journal of Statistical Planning and Inference* **97**, 263–274. [MR1861153](#)
- Johnstone, I. M. and Silverman, B. W. (2005). Empirical Bayes selection of wavelet thresholds. *The Annals of Statistics* **33**, 1700–1752. [MR2166560](#)
- Mallat, S. G. (1989). A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11**, 674–693.
- Morris, J. S. and Carroll, R. J. (2006). Wavelet-based functional mixed models. *Journal of the Royal Statistical Society, Series B* **68**, 179–199. [MR2188981](#)
- Nason, G. P. and Silverman, B. W. (1994). The discrete wavelet transform in  $S$ . *Journal of Computational and Graphical Statistics* **3**, 163–191.
- Percival, D. B. and Walden, A. T. (2000). *Wavelet Methods for Time Series Analysis*. Cambridge: Cambridge Univ. Press. [MR1770693](#)
- Strang, G. (1993). Wavelet transforms versus Fourier transforms. *Bulletin of the American Mathematical Society* **28**, 288–305. [MR1191480](#)

Ecole Polytechnique Fédérale de Lausanne  
EPFL-FSB-MATHAA-STAT  
Station 8, 1015 Lausanne  
Switzerland  
E-mail: [Anthony.Davison@epfl.ch](mailto:Anthony.Davison@epfl.ch)